

**Before the  
Federal Communications Commission  
Washington, D.C. 20554**

<b>In the Matter of</b>	)	
	)	
<b>Digital Broadcast Copy Protection</b>	)	<b>MB Docket No. 02-230</b>
	)	
	)	

**REPLY COMMENTS OF  
CORPORATION FOR NATIONAL RESEARCH INITIATIVES**

Patrice A. Lyons  
Law Offices of Patrice Lyons, Chartered  
910 17<sup>th</sup> St., N.W., Suite 800  
Washington, D.C. 20006  
E-mail: [palyons@bellatlantic.net](mailto:palyons@bellatlantic.net)

Counsel to Corporation for National  
Research Initiatives

February 19, 2003

An important policy objective in a proceeding such as this should be to accommodate, to the greatest extent possible, both the technological variations among current broadcast transmission methods and the technological advancements that are sure to accompany future development. Moreover, it is important to account for the fact that broadcast transmission technologies are being developed within a broader digital information networking environment. Building bridges to the future should be an important consideration in any decision with respect to security for broadcast programming whether made available in digital form or mapped into analog signals for dissemination to the public. While the current proceeding is focused on digital broadcast television, this is just one aspect of a larger information management equation.

Over the last decade, there have been efforts underway to shape a communications environment that will be friendly to both consumers and producers of information. Whether information is made available by satellite or cable, or labeled telephone, cable, wireless or broadcast television, the source may not be evident to the consumer. There is a need to continue the dialogue among the interested parties before taking steps to lock in one or more approaches. A basic starting point would be an identifier system that will thread the various pathways and provide a coherent fabric on which owners of rights in broadcast and other forms of programming may rely in their businesses.

Corporation for National Research Initiatives (CNRI) is a non-profit 501(c)(3) organization that undertakes, fosters, and promotes research in the public interest. Activities center around strategic development of network-based information technologies, and providing leadership and funding for research and development of the National Information Infrastructure. CNRI engages in system and technology research, development and demonstration projects in order to further the design and implementation of selected infrastructure components for new computing- and communications-based applications.

Concerns about misuse of information in connection with broadcast television, whether digital or not, are well founded. In the digital world, these concerns are magnified by the ease with which abuses may occur, since exact replicas of information in digital form can be stored, processed, and disseminated at little or no cost. CNRI has been involved with this issue for many years and has pioneered technology at the infrastructure level that may be helpful in dealing with this matter.

The Handle System<sup>®</sup> is an example of a technology developed by CNRI that may be directly relevant here. It is a highly responsive indirection system on the Internet that resolves unique identifiers (which CNRI calls handles) to “handle records” supplied by the creator or originator of the information being identified. These handle records could contain any usage conditions deemed appropriate for the information. An overview of the current Handle System is contained in Appendix A (additional information is available on the Internet at [www.handle.net](http://www.handle.net)).

If the information being broadcast were to include the relevant usage information from the handle record (it being assumed that the broadcaster would have access to the Internet for the purpose of retrieving and incorporating the usage information in the broadcast program), the device receiving the information could be instrumented to act on that usage information in the specified manner. The end user device would not be required to be connected to the Internet. It would not be required that every device have specific hardware to protect the information from redistribution, but rather that embedded software in the equipment would be required to act appropriately on receipt of the usage information.

For information that is made available on the Internet, similar requirements would apply. The received information would contain the unique identifier, but not necessarily the broadcast flag, and the receiver would be required to resolve the identifier (i.e., handle) to determine the acceptable usage restrictions and to abide by them. This need not be

implemented exclusively in hardware; it could be more flexibly implemented and upgraded in software; and this software need not be directly accessible to the end-user.

The usage information in the handle system could be the “broadcast flag” or such other information as may be desirable or required in the future. Use of the Handle System provides a degree of flexibility that is inherent in indirection systems. These are systems whose main purpose is typically to map between one parameter (such as an identifier or file name) and another (such as an Internet address or a location on a local disk). The computer industry has used this technique for many years, as has the Internet. It is an appropriate mechanism for consideration by the FCC due to its inherent flexibility, and ease of evolution (perhaps using embedded software), while still offering the ability to effectively control usage such as redistribution.

Clearly, this approach requires a buy-in by the relevant parties including content creators, broadcasters and equipment manufacturers. If this were to occur, it is likely that network equipment manufacturers and Internet Service Providers would find the approach palatable as well. Thus, it would appear that this reflects a promising avenue of consideration for all parties.

The Handle System is in widespread use by book publishers who have branded the term Digital Object Identifier (DOI) as their version of a handle. DOIs are handles that begin with the number “10” and resolve to information such as the location of books and electronic journals on the net. The International DOI Foundation (IDF) was established as a membership organization to administer DOIs and the policies and procedures associated with their use.

Use of the proposed broadcast flag approach, while less flexible, has all the inherent drawbacks of any system that can be circumvented. While there may be a requirement to adhere to the restrictions of the broadcast flag, even if the content is not encrypted, all it

takes is one outlaw party to access the relevant bits and make them available. It may be extremely difficult to identify such a party once the damage is done. Peer-to-peer systems have emerged that are difficult to control; and bandwidth limitations that have made illegal sharing of video programming more limited than, say, illegal sharing of songs, will not be with us for very long. These issues need to be addressed now.

Further, the “personal digital network environment” (i.e., home or local environment) is increasingly likely to have wireless components such as one of the recently introduced 802.11 standards. These local networks are accessible from within and without the local environment (although from nearby or short ranges). Although access point controls are available, they can easily be circumvented or not used by the local network, whether in the home or in other areas. Thus, limitations on redistribution using the broadcast flag may be insufficient in such cases. A more detailed discussion on the use of identifiers in the wireless networking world is contained in Appendix B.

The above example drawn from the wireless networking environment is one of many scenarios that should be considered before reaching a decision on how best to manage broadcast programming in digital form. While it may be convenient to adopt a single approach to security for digital broadcast transmission, if past experience is a guide, there is a need to enable a variety of methods to be deployed. Simple measures to encourage interaction among the various interested parties that are drawn from the technological community as well as the creators and disseminators of broadcast programming would appear to be advisable. Agreement on a system of program identification is a logical first step; and most identification systems may be used in connection with the flexible, higher-level handle system.

The use of an indirection mechanism such as the handle system does not relieve the end user of the need to know the law and to abide by it. Thus, any interloper on a local network would be obligated to determine the usage limitations of any material obtained

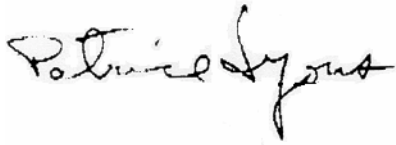
in this fashion. But by using a system of unique, persistent identifiers, and a resolution mechanism such as the handle system, the ability to do so would be accessible from the identifier included with the content itself.

In 1997, a diversified group known as the Cross-Industry Working Team (XIWT) produced an information report entitled "Managing Access to Digital Information: An Approach based on Digital Objects and Stated Operations." A copy of this still timely report, which articulated many of the relevant issues concerning information access in general, and attached hereto as Appendix C, is also available at [www.xiwt.org](http://www.xiwt.org)

Respectfully submitted,

CORPORATION FOR NATIONAL RESEARCH INITIATIVES

By:

A handwritten signature in black ink, appearing to read "Patrice Lyons". The signature is fluid and cursive, with the first name "Patrice" written in a larger, more prominent script than the last name "Lyons".

Patrice A. Lyons  
Law Offices of Patrice Lyons, Chartered  
910 17th St., N.W., Suite 800  
Washington, D.C. 20006  
(202) 293-5990

*Counsel to Corporation for National Research Initiatives*

## **APPENDIX A**

Internet Draft  
Document: draft-sun-handle-system-10.txt  
Expires: March 2003

Sam X. Sun  
CNRI  
Larry Lannom  
CNRI  
September 2002

## Handle System Overview

### Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of Section 10 of RFC2026.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at  
<http://www.ietf.org/ietf/lid-abstracts.txt>  
The list of Internet-Draft Shadow Directories can be accessed at  
<http://www.ietf.org/shadow.html>.

### Abstract

This document provides an overview of the Handle System in terms of its namespace and service architecture, as well as its relationship to other Internet services such as DNS, LDAP/X.500, and URN. The Handle System is a general-purpose global name service that allows secured name resolution and administration over the public Internet. The Handle System manages handles, which are unique names for digital objects and other Internet resources.

### Table of Contents

1. Introduction.....	2
2. Handle Namespace.....	5
3. Handle System Architecture.....	6
4. Handle System Service and its Security.....	9
5. The Handle System and other Internet Services.....	10
5.1 Domain Name Service (DNS).....	11
5.2 Directory Services (X.500/LDAP).....	11
5.3 Uniform Resource Names (URN).....	12



6. Security Considerations.....	13
6.1 General Security Practice.....	13
6.2 Privacy Protection.....	14
6.3 Caching and Proxy.....	14
6.4 Mirroring.....	15
6.5 Denial of Service (DoS).....	15
7. History of the Handle System.....	15
8. Acknowledgement.....	15
References and Bibliography.....	16
Author's Addresses.....	17

## 1. Introduction

This document provides an overview of the Handle System, a distributed information system designed to provide an efficient, extensible, and secured global name service for use on networks such as the Internet. The Handle System includes an open protocol, a namespace, and a reference implementation of the protocol. The protocol enables a distributed computer system to store names, or handles, of digital resources and resolve those handles into the information necessary to locate, access, and otherwise make use of the resources. These associated values can be changed as needed to reflect the current state of the identified resource without changing the handle. This allows the name of the item to persist over changes of location and other current state information. Each handle may have its own administrator(s) and administration can be done in a distributed environment. The Handle System supports secured handle resolution. Security service such as data confidentiality, service integrity, and non-repudiation are provided upon client's request.

The Handle System provides a confederated name service that allows any existing local namespace to join the global handle namespace by obtaining a unique handle system naming authority. Local names and their value-binding(s) remain intact after joining the Handle System. Any handle request to the local namespace may be processed by a service interface speaking the handle system protocol. Combined with the unique naming authority, any local name is guaranteed unique under the global handle namespace.

There are several services that are in use today to provide name service for Internet resources. Among these the Domain Name System (DNS) [2,3] is the most widely used. DNS is designed "to provide a mechanism for naming resources in such a way that the names are mappable into IP addresses and are usable in different hosts, networks, protocol families, internets, and administrative organizations" [3]. The growth of the Internet has raised demands for various extensions to DNS. There are also attempts to use DNS as a general-purpose resource naming system. However, the

importance of DNS in basic network routing has led to great caution in implementing any DNS extension or overloading the DNS for general-purpose resource naming. An additional factor which argues against using DNS as a general-purpose naming service is the DNS administrative model. DNS names are typically managed by the network administrator(s) at the DNS zone level. There is no provision for per-name administrative structure and no facilities for anyone other than the network administrator to create or manage DNS names. This is appropriate for domain name administration but less so for general-purpose resource naming.

The Handle System has been designed from the start to serve as a general-purpose naming service. It is designed to accommodate very large numbers of entities and to allow distributed administration over the public Internet. The handle system data model allows access control to be defined at the level of each handle data. Each handle can further define its own set of administrators that are independent from the network or host administrator.

Traditional URLs (Uniform Resource Locators) [4] allow certain Internet resources to be named as a combination of a DNS name and local name. The local name may be a local file path, or a reference to some local service (e.g. a cgi-bin script). This combination of DNS name and local name provides a flexible administrative model for naming and managing individual Internet resources. However, the URL practice also has some key limitations. Most URL schemes (e.g., http) are defined for resolution only. Any URL administration has to be done either at the local host, or via some other network service such as NFS. Using a URL as a name typically ties the Internet resource to its current network location. For example, a URL will be tied to its local file path when the file path is part of the URL. When the resource moves from one location to another for whatever reason, the URL breaks.

The Handle System is designed to overcome these limitations and to add significant functionality. Specifically, the Handle System is designed with the following objectives:

- . Uniqueness: Every handle is globally unique within the Handle System.
- . Persistence: A handle is not derived in any way from the entity that it names, but is assigned to it independently. While an existing name, or even a mnemonic, may be included in a handle for convenience, the only operational connection between a handle and the entity it names is maintained within the Handle System. This of course does not guarantee persistence, which is a function of administrative care. But it does allow the same name to persist over changes of

location, ownership, and other state conditions. For example, when a named resource moves from one location to another, the handle may be kept valid by updating its value in the Handle System to reflect the new location.

- . Multiple Instances: A single handle can refer to multiple instances of a resource, at different and possibly changing locations in a network. Applications can take advantage of this to increase performance and reliability. For example, a network service may define multiple entry points for its service with a single handle so as to distribute the service load.
- . Extensible Namespace: Existing local namespaces may join the handle namespace by acquiring a unique handle naming authority. This allows local namespaces to be introduced into a global context while avoiding conflict with existing namespaces. Use of naming authorities also allows delegation of service, both resolution and administration, to a local handle service.
- . International Support: The handle namespace is based on Unicode 3.0 [1], which includes most of the characters currently used around the world. This allows handles to be used in any native environment. The handle protocol mandates UTF-8 [5] as the encoding used for handles.
- . Distributed Service Model: The Handle System defines a hierarchical service model such that any local handle namespace may be serviced either by a corresponding local handle service or by the global service or by both. The global service, known as the Global Handle Registry, can be used to dispatch any handle service request to the responsible local handle service. The distributed service model allows replication of any given service into multiple service sites and each service site may further distribute its service into a cluster of individual servers. (Note that local here refers only to namespace and administrative concerns. A local handle service could in fact have many service sites distributed across the Internet.)
- . Secured Name Service: The handle system allows secured name resolution and administration over the public Internet. The handle system protocol defines standard mechanisms for both client and server authentication, as well as service authorization. It also provides security options to assure service integrity and data confidentiality.

- . Distributed Administration Service: Each handle may define its own administrator(s) or administrator group(s). Ownership of each handle is defined in terms of its administrator or administrator groups. This, combined with the handle system authentication protocol, allows any handle to be managed securely over the public network by its administrator at any network location.
- . Efficient Resolution Service: The handle protocol is designed to allow highly efficient name resolution performance. To avoid resolution being affected by computationally costly administration service, separate service interfaces (i.e., server processes and their associated communication ports) for handle name resolution and administration may be defined by any handle service.

This document provides an overview of the handle namespace and service architecture. It also compares the Handle System with other existing Internet services, protocols, and specifications (e.g., DNS [2, 3], URLs [4], X.500/LDAP [6,7,8], and URN [9,10]). Details of the handle system data and service model, as well as its communication protocol, are specified in separate documents. They can be found under the handle system website at <http://www.handle.net>.

## 2. Handle Namespace

Every handle consists of two parts: its naming authority, otherwise known as its prefix, and a unique local name under the naming authority, otherwise known as its suffix:

`<Handle> ::= <Handle Naming Authority> "/" <Handle Local Name>`

The naming authority and local name are separated by the ASCII character "/". The collection of local names under a naming authority defines the local handle namespace for that naming authority. Any local name must be unique under its local namespace. The uniqueness of a naming authority and a local name under that authority ensures that any handle is globally unique within the context of the Handle System.

For example, "10.1045/january99-bearman" is a handle for an article published in D-Lib magazine [12]. Its naming authority is "10.1045" and its local name is "january99-bearman". The handle namespace can be considered as superset of many local namespaces, with each local namespace having a unique naming authority under the Handle System. The naming authority identifies the administrative unit of creation, although not necessarily continuing administration, of the associated handles. Each naming authority is guaranteed to be

globally unique within the Handle System. Any existing local namespace can join the global handle namespace by obtaining a unique naming authority so that any local name under the namespace can be globally referenced as a combination of the naming authority and the local name as shown above.

Naming authorities under the Handle System are defined in a hierarchical fashion resembling a tree structure. Each node and leaf of the tree is given a label that corresponds to a naming authority segment. The parent node presents the parent naming authority of its child nodes. Unlike DNS, handle naming authorities are constructed left to right, concatenating the labels from the root of the tree to the node that represents the naming authority. Each label is separated by the octet used for ASCII character "." (0x2E). For example, a naming authority for the National Digital Library Program ("ndlp") at the Library of Congress ("loc") is defined as "loc.ndlp".

Each naming authority may have many child naming authorities registered underneath. Any child naming authority can only be registered by its parent after its parent naming authority is registered. However, there is no intrinsic administrative relationship between the namespaces represented by the parent and child naming authorities. The parent namespace and its child namespaces may be served by different handle services, and they may or may not share any administration privileges between each other.

Handles may consist of any printable characters from the Universal Character Set (UCS-2) of ISO/IEC 10646, which is the exact character set defined by Unicode v2.0 [1]. The UCS-2 character set encompasses most characters used in every major language written today. To allow compatibility with most of the existing systems and prevent ambiguity among different encoding, the handle system protocol mandates UTF-8 to be the only encoding used for handles. The UTF-8 encoding preserves any ASCII encoded names so as to allow maximum compatibility to existing systems without causing naming conflict. Some encoding issues over the global namespace and the choice of UTF-8 encoding are discussed in [13].

By default, handles are case sensitive. However, a handle service may define its namespace so that ASCII characters within any handle under the namespace are case insensitive.

### 3. Handle System Architecture

The Handle System defines a hierarchical service model. The top level consists of a single global service, known as the Global Handle Registry (GHR). The lower level consists of all other handle services, generically known as Local Handle Services (LHS).

The Global Handle Registry can be used to manage any handle namespace. It is unique from any other handle services only in that it provides the service used to manage naming authorities, all of which are managed as handles. The naming authority handle provides information that clients can use to access and utilize the local handle service for handles under the naming authority.

Local Handle Services are intended to be hosted by organizations with administrative responsibility for handles under certain naming authority. A Local Handle Service may be responsible for any number of local handle namespaces, each of which identified by a unique naming authority. The Local Handle Service and its responsible set of local handle namespaces must be registered under the Global Handle Registry.

One important aspect of the Handle System is its distributed architecture. The Handle System as a whole consists of a number of individual handle services. Each of these service may consist of one or more service sites. Each of these service site is a complete replication of each other, at least for handle resolution. Additionally, a service site may also consist of one or more handle servers. Handle requests directed at the service site may be evenly distributed into these handle servers. The Handle System may consist of any number of handle services. There are no design limits on the number of sites which make up each service. Neither there are any limits on the number of servers that make up each site. Replication among any service sites does not require that each site contains the same number of servers. In other words, while each site will have the same replicated set of handles, each site may allocate that set of handles across a different number of servers. This distributed approach is intended to aid scalability to accommodate any large-scale of operation and to mitigate problems of single point failure.

Figure 3.1 illustrates a potential handle service that consists of two service sites: one located at the US East coast and the other at the US West coast. The East coast service site consists of four server computers. The West coast service site, with more powerful computers deployed, decides two servers will suffice. The number of service sites for any handle service, as well as the number of servers that are used by any service site, may be added or removed dynamically depending on the service requirement.

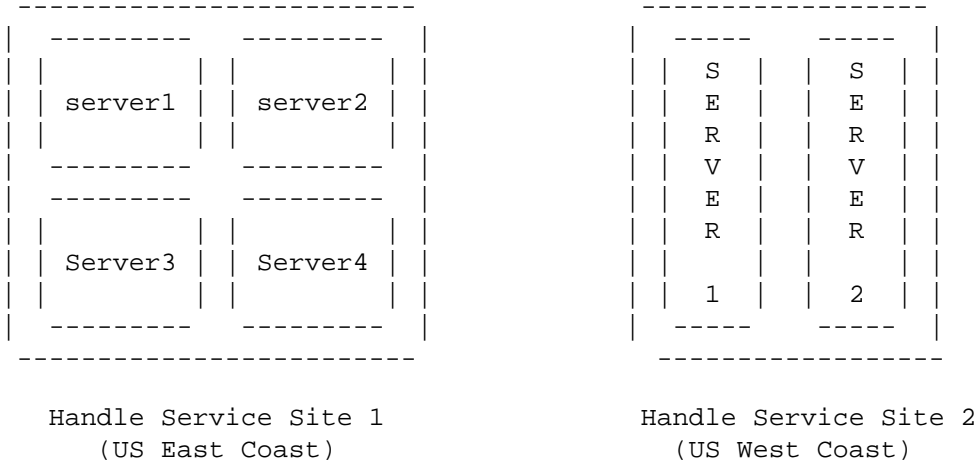


Fig. 3.1: Handle service configured with two service sites

Each handle service manages a distinct sub-namespace under the Handle System. Namespaces under different handle services may not overlap. The sub-namespace typically consists of handles under a number of naming authorities. The handle service is called the "home" service of these naming authorities and is the only one that provides resolution and administration service for handles under these naming authorities. Before resolving a handle, a client has to determine the "home" service of the handle in question. The "home" service of each handle is the "home" service of its naming authority and is registered at the Global Handle Registry. Clients can find the "home" service for each handle by querying the naming authority handle at the Global Handle Registry.

The Global Handle Registry maintains naming authority handles. Each naming authority handle maintains the service information that describes the "home" service of the naming authority. The service information lists the service sites of the handle service, as well as the interface to each handle server within each site. To find the "home" service for any handle, a client can query the Global Handle Registry for the service information associated to the corresponding naming authority handle. The service information provides the necessary information for clients to communicate with the "home" service.

Figure 3.2 shows an example of a typical handle resolution process. In this case, the "home" service is a Local Handle Service. The client is trying to resolve the handle "cnri.dlib/july95-arms" and has to find its "home" service from the Global Handle Registry. The "home" service can be found by sending a query to the Global Handle Registry for the naming authority handle for "cnri.lib". The Global Handle Registry returns the service information of the Local Handle

Service that is responsible for handles under the naming authority "cnri.dlib". The service information allows the client to communicate with the Local Handle Service to resolve the handle "cnri.dlib/july95-arms".

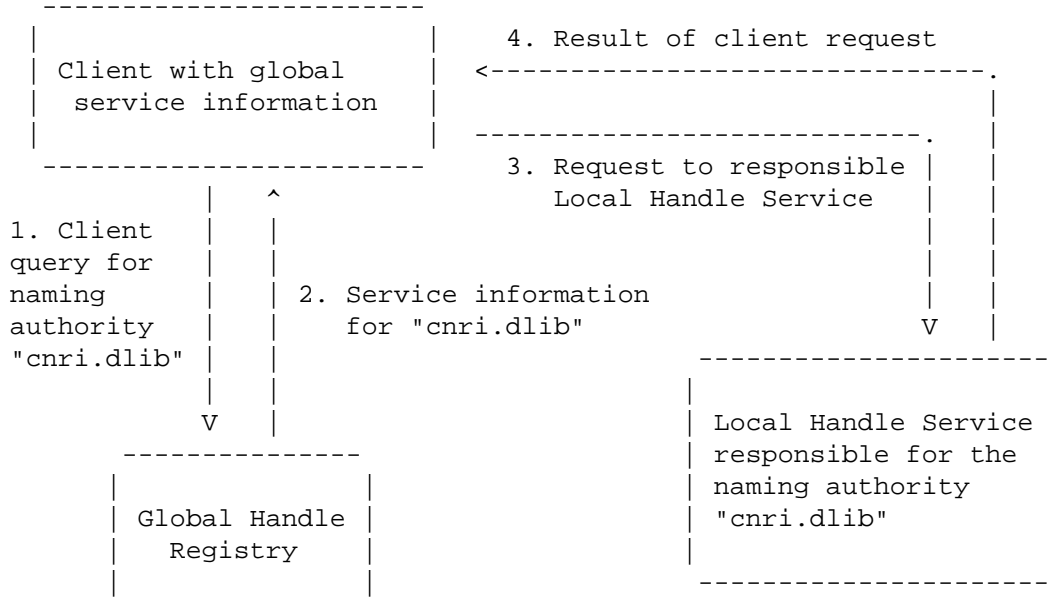


Fig. 3.2: Handle resolution starting with global

To improve resolution performance, any client may choose to cache the service information returned from the Global Handle Registry and use it for subsequent queries. A separate handle caching server, either stand-alone or as a piece of a general caching mechanism, may also be used to provide shared caching within a local community. Given a cached resolution result, subsequent queries of the same handle may be answered locally without contacting any handle service. Given cached service information, clients can send their requests directly to the Local Handle Service without contacting the Global Handle Registry.

#### 4. Handle System Service and its Security

The Handle System provides handle resolution and administration service over the public Internet. Each handle can be assigned with a set of values. Clients use the handle resolution service to resolve any handle into its set of values. Each value has a data type and a unique value index. Clients can query for specific handle values based on data type or value index.



The handle administration service answers requests from clients to manage handles. These include adding handles, deleting handles or updating their values. It also manages naming authorities via naming authority handles. Each handle can define its own administrator(s) and each administrator can be granted a certain set of permissions. The handle system authentication protocol authenticates the handle administrator before fulfilling any administrative request.

The Handle System provides security services such as client and server authentication, data confidentiality and integrity, as well as service non-repudiation. By default, handle resolution service does not require any client authentication. However, resolution request for confidential data assigned to any handle (by its administrator), as well as any administration request (e.g. adding or deleting handle values) require authentication of the client for proper authorization. When authentication is required, the handle server will issue a challenge to the requesting client before carrying out the client's request. To satisfy the authentication requirement, the client must send back the correct response that identifies itself as the administrator. The handle server will respond to the initial request only after successful authentication of the client. Handle clients may choose to use either secret key or public key cryptography for authentication. Authentication under Handle System can also be carried out via third party authentication services. To ensure data integrity, clients may request digitally signed responses from any handle server. They may also set up a secured communication session with the handle server so that any exchanged information can be encrypted (for data confidentiality) using the session key.

The Handle System provides service options for secured information exchange between client and server. This does not guarantee the truthfulness of handle values. Incorrect values assigned to any handle by its administrator may very well mislead clients. On the other hand, a handle value may contain references to other handle values to provide additional credentials. For example, a handle value R (e.g., a claim) may contain a reference to some other handle value that contains the digital signature (from a creditable source) upon the value R. Clients who trust the signature could then trust the handle value R.

## 5. The Handle System and other Internet Services

There are a number of existing and proposed Internet identifier services or specifications that by design or intent cover some of the functionalities proposed for the Handle System. This section briefly reviews them in relationship to the Handle System.

## 5.1 Domain Name Service (DNS)

The Domain Name Service, or DNS, was originally designed and is heavily used for mapping domain names into IP Addresses for network routing purposes. RFC1034 [2] and RFC1035 [3] provide detailed descriptions of its design and implementation. The growth of the Internet has increased demands for various extensions to DNS, even its possible use as a general purpose resource naming system. However, any such use has the potential to slow down the network address translation and/or affect its effectiveness in network routing. DNS implementations typically do not scale well when large amount of data is associated with any particular DNS name. It is generally considered inadequate to use DNS for naming any kind of resources over the Internet.

An additional factor that argues against using DNS as a general-purpose naming service is the DNS administrative model. DNS names are typically managed by the network administrator(s) at the DNS zone level. There is with no provision for a per-name administrative structure. No facilities are provided for anyone other than network administrators to create or manage DNS names. This is appropriate for domain name administration but less so for general-purpose name administration.

The Handle System differs from DNS in its distributed administration and service model, as well as its security features. The handle system protocol comprise security options to assure confidentiality and integrity during data transmission. Each handle under the Handle System may define its own administrator that is independent from the server administrator. The handle system protocol allows any handle administrator to manage its handles securely over the public network. Additionally, the Handle System service model allows any of its service sites to dynamically configure its service distribution among a cluster of servers to accommodate increased service requests. This also allows less powerful computers to be used together to support any huge number of handles.

## 5.2 Directory Services (X.500/LDAP)

X.500 [6] is the OSI Directory Standard defined by ISO and the ITU. It is designed "to provide a white pages service that would return either the telephone numbers or X.400 O/R addresses of people", and is "concerned mainly with providing the name server service for Open Systems Interconnection (OSI) applications" [7]. X.500 defines a hierarchical data and information model with a set of protocols to allow global name lookup and search. The protocol, however, has proved difficult to implement and there has been difficulty in

getting "client access integrated into existing products" [14]. LDAP (Lightweight Directory Access Protocol) [8] has overcome many of these difficulties by making the protocol simpler, and easier to implement. Some concern remains, however, that as LDAP is emerging from a local directory access protocol (LDAP v2) into a distributed service protocol (LDAP v3), it faces many issues not addressed in its original design, resulting in new complications.

The fundamental difference between a name resolution service such as the Handle System and a directory service such as LDAP is search capability. The added functionality of being able to search a directory service necessarily carries with it added complexity, thus affects its efficiency. A pure name service, such as the Handle System, can be designed solely around efficient resolution of known items without addressing functions and data structures required for discovery of unknown items based on incomplete criteria.

Directory services such as LDAP or WHOIS++ [15,16] may be used in tandem with the Handle System to provide reverse lookup service. Existing corporate directory services, for example, could provide interfaces to both services. The handle system interface would provide a highly efficient name resolution service. The directory service interface would provide extended search capability. Handles could also be used in LDAP service referral. For example, a LDAP service may be referenced as a handle. Doing so will make the reference persistent overtime, independent from location change.

### 5.3 Uniform Resource Names (URN)

The IETF URN Working Group [11] has defined a syntax, possible resolution mechanisms, and namespace registration procedure for a resource identifier intended to cover a large array of existing and potential namespaces. Namespaces are to be registered and assigned unique Namespace Ids (NIDs). Any resolution services associated with these namespaces require further registration with a Resolution Discovery System (RDS) which clients could use to begin, or discover, the appropriate resolution mechanisms.

The objectives and some of the approaches of the URN and Handle System efforts have enough in common that some observers might think that they are in contention. This is not the case. The URN effort is explicitly designed to accommodate multiple identifier namespaces and resolution systems. The Handle System is one such case. It has a very specific data and service model, along with a protocol that supports both handle resolution and administration. URNs and the Handle System may interact in variety of ways. The most obvious of which is that the Handle System could be registered as a URN namespace. In other words, handles under the Handle System

could be referenced as a type of URN. On the other hand, it would also be possible to use the Handle System as a type of RDS for other URN namespaces. The success of either system however, is not dependent upon the success of the other.

## 6. Security Considerations

This section is meant to inform people of security limitations of the Handle System, as well as precautions that should be taken by application developers, service providers, and handle system clients. Specific security considerations regarding the handle system protocol [21] or its data and service model [22] are addressed in separate documents.

### 6.1 General Security Practice

The security of the Handle System depends on both client and server host security at every step in the transaction. It assumes the client host has not been tampered with and that client software will reliably convey the received data to the client. The client of any handle service must also assume that any handle servers involved have not been compromised. To trust the Global Handle Registry is to believe that the Global Handle Registry will rightfully direct the client request to the responsible Local Handle Service. To trust a Local Handle Service is to believe that the Local Handle Service will correctly return the data that was assigned to the handle by its administrator. A Local Handle Service typically supports a set of naming authorities. Thus, trusting a Local Handle Service would imply trusting those naming authorities.

The handle system service integrity depends heavily on the integrity of the global service information. Invalid global service information may mislead clients into inappropriate Local Handle Services. It may also allow attackers to forge server signatures. The Global Handle Registry must take extreme caution in protecting the global service information and the public key pair used to sign the global service information. Client applications should only accept the global service information from the Global Handle Registry. They should check its integrity upon each update.

For efficiency reasons, handle servers will not generate or return digital signature for every service response unless specifically requested by clients. To assure data integrity, clients must explicitly ask the server to return the digital signature. To protect sensitive data from exposure, clients may establish a communication session with the server and ask the server to encrypt any data using the session key.

## 6.2 Privacy Protection

By default, most handle data stored in the Handle System is publicly accessible unless otherwise specified by the handle administrator. Handle administrators must pay attention when adding handle values that contain private information. They may choose to mark these handle values readable only by the handle administrator(s), or to store these handle values encrypted, so that these values can only be readable within a controlled set of audience.

Log files generated by the handle server are another vulnerable point where client privacy may be under attack. Operators of handle servers must protect such information carefully.

## 6.3 Caching and Proxy

Besides performance gains and other value-added services, both the proxy and caching server present themselves as men-in-the-middle, and as such are vulnerable to man-in-the-middle attacks. It is important to know that proxy and caching servers are not part of any handle service. They are clients of the Handle System. Service responses from proxy and caching servers cannot be authenticated via handle system protocol. The trust between the client and its proxy and caching server has to be setup independently.

By using the proxy and caching server, clients assume that the server will submit their request and relay any response from the Handle System, without mishandling any of the contents. They also assume that the server will protect any sensitive information on their behalf.

Proxy and caching server operators should protect the systems on which such servers are running as they would protect any system that contains or transports sensitive information. In particular, log information gathered at proxies often contain highly sensitive personal information, and/or information about organizations. Such information should be carefully guarded, and appropriate guidelines for their use developed and followed.

Caching servers provide additional potential vulnerabilities because the contents of the cache represents an attractive target for malicious exploitation. Potential attacks on the cache can reveal private data for a handle user, or information still kept after a user believes that they have been removed from the network. Therefore, cache contents should be protected as sensitive information.

## 6.4 Mirroring

Handle system clients should be aware of possible delays in content replication among mirroring sites. They should consider sending their request to the primary service site for any time-sensitive data. Selection of mirroring sites by service administrator must be done carefully. Each mirroring site must follow the same security procedures in order to ensure the service integrity. Software tools may be applied to ensure data consistency among mirroring sites.

## 6.5 Denial of Service (DoS)

As with any public service, the Handle System is subject to denial of service attack. No general solutions are available to protect against such attack in today's technology. Server implementations may be developed to be aware of such attack and notify its administrator when it happens. Stateless cookies [19, 20] are one means to mitigate some of the effects of DoS attacks on hosts that perform authentication, integrity, and encryption services. Server implementations, moreover, need to be upgradeable to take advantage of new security technologies including anti-DoS technologies as these become available.

## 7. History of the Handle System

The Handle System was originally conceived and developed at CNRI as part of the Computer Science Technical Reports (CSTR) project, funded by the Defense Advanced Projects Agency (DARPA) under Grant Number MDA-972-92-J-1029. One aspect of this early digital library project, which was also a major factor in the evolution of the Networked Computer Science Technical Reference Library (NCSTRL) [18] and related activities, was to develop a framework for the underlying infrastructure of digital libraries. It is described in a paper by Robert Kahn and Robert Wilensky [17]. The first implementation was created at CNRI in the fall of 1994 in an effort led by David Ely.

Early adopters of the Handle System include the Library of Congress, the Defense Technical Information Center (DTIC), and the International DOI Foundation (IDF). Feedback from these organizations as well as NCSTRL, other digital library projects, and related IETF efforts as mentioned above have all contributed to the evolution of the Handle System. Current status and available software, both client and server, can be found at <http://www.handle.net>.

## 8. Acknowledgement

This work is derived from the earlier versions of the handle system implementation. Design ideas are based on those discussed within the handle system development team, including David Ely, Charles Orth, Allison Yu, Sean Reilly, Jane Euler, Catherine Rey, Stephanie Nguyen, Jason Petrone, and Helen She. Their contributions to this work are gratefully acknowledged.

The authors also thanks and acknowledges Mark Baugher ([mbaugher@cisco.com](mailto:mbaugher@cisco.com)) for his extensive review and comments of these drafts, as well as recommendations received from other members of the IRTF IDRM research group (<http://www.idrm.org>).

## References and Bibliography

- [1] The Unicode Consortium, "The Unicode Standard, Version v3.0", Addison-Wesley Pub Co; ISBN: 0201616335
- [2] P. Mockapetris, "DOMAIN NAMES - CONCEPTS AND FACILITIES", RFC1034, November 1987
- [3] P. Mockapetris, "DOMAIN NAMES - IMPLEMENTATION AND SPECIFICATION", RFC1035, November 1987
- [4] Berners-Lee, T., Masinter, L., McCahill, M., et al., "Uniform Resource Locators (URL)", RFC1738, December 1994
- [5] Yergeau, Francois, "UTF-8, A Transform Format for Unicode and ISO10646", RFC2044, October 1996
- [6] ITU-T Rec. X.500, "The Directory: Overview of Concepts, Models, and Services", 1993.
- [7] D W Chadwick, "Understanding X.500 - The Directory", Chapman & Hall ISBN: 0-412-43020-7.
- [8] Wahl, M., Howes, T., and S. Kille, "Lightweight Directory Access Protocol (v3)", RFC 2251, December 1997
- [9] Sollins, K., and L. Masinter, "Functional Requirements for Uniform Resource Names", RFC 1737, December 1994
- [10] Sollins, K. "Architectural Principles of Uniform Resource Name Resolution", RFC 2276, January 1998
- [11] IETF Uniform Resource Names (URN) Working Group, April, 1998, <http://www.ietf.org/html.charters/urn-charter.html>
- [12] D-Lib Magazine, <http://www.dlib.org>

- [13] Sam X. Sun, "Internationalization of the Handle System - A Persistent Global Name Service", Proceeding of 12th International Unicode Conference, April, 1998
- [14] D Goodman, C Robbins, "Understanding LDAP & X.500", August 1997
- [15] Deutsch P., Schoultz R., Faltstrom P., and C. Weider, "Architecture of the Whois++ service", RFC 1835, August 1995
- [16] Weider, C., J. Fullton, and S. Spero, "Architecture of the Whois++ Index Service", RFC 1913, February 1996
- [17] Kahn, Robert and Wilensky, Robert. "A Framework for Distributed Digital Object Services", May, 1995
- [18] The Networked Computer Science Technical Reports Library (NCSTRL), <http://www.ncstrl.org/>
- [19] P. Karn, W. Simpson, "Photuris: Session-Key Management Protocol", March, 1999
- [20] D. Harkins, D Carrel, "The Internet Key Exchange (IKE)", November, 1998
- [21] S. Sun, S. Reilly, L. Lannom, "Handle System Namespace and Service Definition", IETF draft, <http://www.ietf.org/internet-drafts/draft-sun-handle-system-def-05.txt>, work in progress.
- [22] S. Sun, S. Reilly, L. Lannom, J. Petrone, "Handle System Protocol Specification", IETF draft, <http://www.ietf.org/internet-drafts/draft-sun-handle-system-protocol-02.txt>, work in progress.

#### Author's Addresses

Sam X. Sun  
Corporation for National Research Initiatives (CNRI)  
1895 Preston White Dr. Suite 100  
Reston, VA 20191  
Phone: 703-262-5316  
Email: [ssun@cnri.reston.va.us](mailto:ssun@cnri.reston.va.us)

Larry Lannom  
Corporation for National Research Initiatives (CNRI)  
1895 Preston White Dr. Suite 100  
Reston, VA 20191  
Phone: 703-620-8990  
Email: [llannom@cnri.reston.va.us](mailto:llannom@cnri.reston.va.us)



## **APPENDIX B**

## **USE OF IDENTIFIERS IN THE WIRELESS NETWORKING WORLD**

It is likely that broadcast programming will be made available to wireless devices in various digital formats that comply with the requirements of IEEE standard 802.11. What happens to a program when it is accessed by a user may not resemble the current over the air technology. Any identifier such as the proposed “broadcast flag” should be flexible enough to accommodate such developments. To do otherwise would constrain the creativity of information providers as they experiment with new forms of expression. There is also the vulnerability that may be experienced where information is not managed in a more coordinated manner along the communications pathways from source to user. Many different players will be called upon to coordinate their technologies and business plans.

A starting point in any analysis of a system of content identification and, more generally, information management with respect to wireless networking technology is a clear understanding of what is meant by “content” in this context. This has important ramifications for intellectual property owners and information systems developers alike. Because of the broader implications of the development of such systems, an effort must be made at the outset to allow them to be compatible or interoperable at all levels. For example, the use of higher level identifiers to authorize or coordinate simple lower level functions such as “change status to activate or not” may assist in the provision of digital information and other digital goods and services that may be subject to intellectual property restrictions; and such identifiers may also be required for the management of more complex types of wireless networking operations.

Despite the unfortunate tendency to view “content” as restricted to traditional copyright works, whether “born digital” or converted to digital form, a much wider variety of digital “content” is being implemented in a network environment. Broadcast

programming may evolve to encompass a wide variety of information. When organizing and deploying identifier systems at lower 802.11 levels, care should be taken to accommodate the interaction between 802.11 identifier systems and new forms of “content.” Examples of such content might take the form of virtual machines implemented at various points in an 802.11 network, as well as the static and dynamic relationships of such information management systems or other identifiable elements in the Internet or other networking environment.

Much effort has been expended on the development of various computational facilities known as “virtual machines;” and these facilities are being used to perform various operations for a wide variety of home and business applications. Where such virtual machines interact with information elements in an underlying 802.11 physical layer, manage access to external digital resources or perform other “stated operations,” these computational facilities may themselves be viewed as “content.” The interaction of identifiers assigned to this new form of “content” at various levels of granularity, as well as the identifiers used for more traditional copyright resources, and the interaction of any such identifiers with lower level 802.11 information elements, is an important area for further consideration within the IEEE 802.11 standards process.

Where content providers have developed digital asset management systems to identify their digital goods and services, including specialized metadata and related rights management technology, the tracking of such goods and services may be important for owners of intellectual property rights. Several concepts used in 802.11 may require reassessment to accommodate this development.

While the issue of digital asset management might be seen as unimportant to the developers of wireless networks, whose focus may only be on communications connectivity, it could be of real concern to owners of content, including broadcast stations, who may have no other effective recourse to monitor or verify compliance with

terms and conditions placed on specific digital information goods or services. The ability for a sufficiently endowed (i.e., with powerful computer resources) unauthorized “outside party” to tap into an 802.11 network, effectively subverting normal security provisions, is clearly within the realm of possibility. If such concerns become prevalent, it may be desirable to assign some form of identifier external to the 802.11 specification to each transmitted frame, or specific elements of a frame, so that the identifier can be used to maintain records of authorized transactions or to track any unauthorized use or interception back to its source (indeed, a “frame” may itself be viewed as a structured information resource in this context, i.e., content with its own identifier and stated operations).

A question may also be raised with respect to other information elements and associated element identifiers set forth in 802.11 (*IEEE 802.11 Handbook* by Bob O’Hara and Al Petrick (1999), at page 67). Any identifiers or other metadata associated with “access points,” “stations,” “MACs,” “frames,” or other 802.11 compliant elements, could be made known to digital information management systems, or mutually trusted third parties, and steps taken to coordinate such 802.11 identifiers with identifiers associated with “content,” whether or not such identifiers are external to the 802.11 standard. This would appear to be a useful step toward encouraging the development of commerce based on wireless networking technology where intellectual property, security, privacy or other restrictions apply.

At the present time, virtually all computer-based communication systems involve moving bits from a source to one or more destinations (in the latter case this may occur by broadcast or selective multicast as well as multiple one-to-one interactions) without regard to the meaning of the bits being communicated. For purposes of content identification, it would be most useful and practical to identify content at higher levels than either 802.3 or 802.11 now appears to allow. For example, if content capable of being independently identified and processed was present in the form of a “digital object”

(i.e., structured data having an associated unique persistent identifier), then it would be possible to track content at various points in the communications pathway, or even identify transaction records at such points. However, since actual content may be encrypted in different networks or, more generally, in different information systems, explicit arrangements would have to be made with the system operators to leave the identifier field in the clear (if, indeed, the 802.11 standard would allow this when encryption is used), or to trust various intermediary systems along the way that see the content in the clear to extract the identifiers for the purpose of content identification and processing.

At best, this is a very sensitive matter. Any such arrangements would have to be built into agreements with the originators of the content and managed within the overall communications environment. This may be accomplished through such means as associating specific terms and conditions with individual digital objects in a form that is interpretable along the way so that appropriate decisions can be made on the performance of permitted operations such as further dissemination, reproduction or aggregation. An example of an identification/ resolution system that can assist here is the Handle System® (see <http://www.handle.net>).

The assignment and use of identifiers associated with digital objects and other digital resources is an important area of research. Some progress has been made in this context, but much remains to be done. Coordination of these efforts with the IEEE 802.11 standard development process, and related efforts, is desirable. While 802.11 may be viewed by some as too low-level a system to encumber with this kind of baggage, certain basic identification elements might be desirable at that level to authenticate information systems and other digital resources to facilitate verification and compliance with approved “stated operations” for each digital object or other digital resource (whether also viewed as a “MIB,” “communication” or “frame”) without violating any confidences or other restrictions placed on the material.

It is also important to provide a logical distributed connection between any such lower level identifiers with intermediate management system identifiers, e.g., URIs or other names assigned to various elements in one or more virtual machines that may be connected with an 802.11 computational facility, and, ultimately, with identifiers and other metadata that may be associated with information elements by intellectual property owners or their agents for purposes not just of delivery of digital objects or other digital resources, but to enable a wide variety of stated operations to be performed on such objects on a static or dynamic basis.

In summary, at a minimum, there is a need for visibility between higher-level identifiers and those assigned at lower levels such as MAC addresses, as well as the coordination of these identifiers and related metadata with network system elements such as IP addresses. This may take the form of simple methods for tracking and accounting of identifiable digital information in order to facilitate the enforcement of contractual restrictions on material subject to intellectual property, or the detection of unauthorized external intrusions.

Patrice A. Lyons

## **APPENDIX C**

---

***Managing Access to  
Digital Information:  
An Approach Based on Digital  
Objects and Stated Operations***

---

*May 1997*

*3Com  
Alcatel Telecom  
American Management Systems  
Apple Computer  
AT&T  
BBN  
Bell Atlantic  
Bellcore  
BellSouth  
Cisco  
Citicorp  
Compaq  
Corning  
CyberCash  
Digital Equipment  
EarthLink Network  
Electric Power Research Institute  
Ericsson  
Fujitsu  
GTE Laboratories  
Hewlett-Packard  
Houston Associates  
Hughes Network Systems  
IBM  
Intel  
InterTrust  
Lucent Technologies  
MCI Communications  
Motorola  
NEC USA  
New York Times  
NIST  
Nortel (Northern Telecom)  
Novell  
Philips Research Briarcliff  
Prodigy Services  
QuantumLink  
Science Applications International Corporation  
Silicon Graphics  
Southwestern Bell  
Sprint  
Sun Microsystems  
Texas Instruments  
USWest  
West Group*



## ***Table of Contents***

1.0 Introduction

2.0 Digital Objects

3.0 Key Infrastructure Requirements

4.0 Potential Business Opportunities

5.0 Access to Digital Objects

6.0 Reflections on Rights Management Technologies

7.0 Conclusion and Recommendations

References

Appendix A

Appendix B

Note: Permission is hereby granted to reproduce, distribute, and/or disseminate portions of this report solely for non-commercial information purposes, provided that credit is given to XIWT. For other uses, please contact XIWT. © Copyright 1997. All Rights Reserved.

## 1.0 Introduction

---

The deployment and widespread use of global information systems like the Internet has dramatically reduced the production and distribution costs usually associated with information dissemination. With digital technology, anyone can become an information provider, able to generate and distribute ideas at little or no cost. This technology facilitates dynamic and efficient forms of creativity or innovation, such as the integration of various digital materials to manifest sounds, images, graphics, or industrial designs - all linked in interesting, innovative, or entertaining ways. It offers wider, more dynamic forms of collaboration. For example, works of literature, music, or art converted to digital form, or initially expressed in some digital format, can now be worked on by large numbers of collaborators separated by time and space. New works can be produced in record time; and scientific theories or problems can be proposed, decomposed, simulated, and worked on in parallel almost in real time by researchers around the world.

In the past, there was a relative shortage of creative work. A network environment may change this situation. In an open, accessible computer network environment, even the smallest voice can be widely heard. The result will be new dynamics for the economics of content production and distribution. Without the ability to control access to information in a network environment, however, intellectual property may have little value.

Although the future "networked digital world" holds promise for greater societal good, it presents new challenges with respect to existing legal systems. Information expressed in various digital formats is easy to reproduce, perform, and disseminate with nearly perfect accuracy at low cost. Digital information can be made immediately accessible to everyone without regard to location. Further, information technology allows for more interactivity. Links between works can be dynamically made and broken; and composite objects and works can be rapidly composed, nested, and/or transformed almost effortlessly (except for the intellectual energy expended).

Many laws apply to digital information, including the laws of copyright, patent, trademark, libel, slander, defamation, contract, and communications - not to mention the First Amendment to the U.S. Constitution. While copyright, patent, trademark, and communications laws are among the more important bodies of law in this context, there will also be instances where legal provisions in such areas as contracts, taxes, securities, banking, insurance, and trade arrangements come into play. Legal provisions often interact, and, occasionally, overlap or even conflict in practice. Storing, manipulating, accessing, and distributing digital objects and other digital resources, and executing new types of operations on digital material, bring to bear a wide variety of laws and regulations. Thus, while it is important to understand the implications of each of these legal systems individually, as a matter of public policy, it is necessary to consider the combined effects of the various elements in light of new technological developments.

Emerging technical capabilities - especially efforts to develop data structures for use in a digital environment - may in fact advance public understanding of the relevant legal implications. Thus, for example, business models are being developed around the concept of a "container" or "package" that may embody digital information subject to various rights or interests or that, when processed, may manifest such "content." In this paper, we refer to such data structures as **digital objects**. Conceptually, a digital object is a logical entity or data structure whose two principal components are digital material ("data"), plus a unique identifier for the material and other information pertaining to the data ("metadata").

Emerging cryptographic and agent-based technologies may make it easier to manage rights such as copyright by (1) providing more effective enforcement through cryptography, secure hashing/digital fingerprints, and certificates; (2) facilitating more efficient payment through digital cash and micropayments; and (3) providing efficient means for monitoring and detecting infringements through the use of intelligent agents. Cryptography can be used to give each digital object a unique digital signature; combined with steganographic technology, it can be used to imbed hidden or invisible markings or fingerprints. These digital object fingerprints then can be used to test for authenticity. This technology enables digital marks and digital signatures to be placed on digital objects in such a way that they cannot be copied or removed without detection. When misapplied, however, these technologies may discourage and inhibit the very sharing and cross-fertilization they are meant to encourage.

These developments do not diminish the importance of intellectual property protection. In fact, there is a growing need to encourage new forms of authorship or discovery that do not just replicate or mirror old forms, but that reconceptualize what it means to be a creative work or invention. It is important, however, not to lose sight of the cost factors associated with the creative process. While much new information may now be made accessible in a network environment, it may actually offer little in the way of real creativity, inventiveness - or even interest. It will be increasingly important to find ways to reward those who add significant value to the store of human knowledge, who entertain, or who collect and disseminate information, and to encourage them to share their work widely with others. This paper addresses various issues surrounding the management of rights and permissions in the digital environment. It introduces, in particular, the notion of digital objects (sometimes referred to as packages, containers, or structured bit sequences) and their supporting technologies as a means of enabling new business opportunities and protecting intellectual property in a computer network environment.

## 2.0 Digital Objects

---

Digital objects provide a means of organizing and identifying "content" - i.e., underlying data - for purposes of storage, access, or distribution. A digital object is not merely an unstructured sequence of bits or symbols from an alphabet. Rather, it has a structure that

allows it to be identified and its content to be organized and protected, as appropriate. As described by Kahn and Wilensky (1995), a digital object may incorporate (or be interpreted to manifest) not only the content, but also the unique identifier of the digital object and other metadata about the digital object and its content. The metadata may include restrictions on access to digital objects, notices of ownership, and licensing agreements relating to underlying content.

A digital object may also be viewed as information in its own right, with its own intrinsic rules and procedures, and may itself be packaged in other digital objects. The packaging of a digital object within another may occur, for example, where agent software is charged with accessing information on behalf of a user, or collecting and organizing information that it presents to a user. Digital objects may be stored in repositories which, in turn, may be structured as digital objects - i.e., logical entities containing multiple digital objects.

The notion of a digital object as a container that incorporates protected information may facilitate the development of flexible and efficient mechanisms for managing rights and interests in protected information within a network environment. There will probably be at least two different categories of digital objects - those that come with meaningful restrictions and those that do not. Many commercial digital objects may come without any meaningful restrictions; others may be heavily encumbered.

Defining a digital object infrastructure allows business models to be developed that can be based on communications law and other bodies of law. It is anticipated that these legal systems will provide an adequate basis for managing access to digital objects in order to perform "stated operations" and provide related services.

## **2.1 Agents as Digital Objects**

Software agents are a particularly interesting technology for managing rights and executing tasks in the network environment. When configured as digital objects, they may act on behalf of rights-holders to protect works embodied in such objects, and they may interact with other agents and systems to carry out a wide range of tasks in the networked digital world.

Briefly, software agents are computer programs that may be mobile in a network environment and can act as intermediaries providing information about rights and permissions. Agents can be used to control the distribution of material, discover infringements of rights and interests in intellectual property, and negotiate licenses in a network environment. They can be more than mere transport mechanisms for connection purposes. For example, they can also combine, filter, index, rearrange, interpret, and transform digital information. They could serve as a researcher's assistant - reading the works of others, and then rearranging and reinterpreting them, rather than merely reporting and regurgitating the works verbatim.

## **2.2 Some Issues for Further Exploration**

A variety of questions arise with respect to digital objects and their contents in a network environment. These fall into four general categories - incentive issues, legal issues, business issues, and technology issues. A sampling of the key questions within each of these categories follows.

### **2.2.1 Incentive Issues**

There are several questions that need to be addressed to make sure that producers, consumers, and network service providers are comfortable with this new means of information access. As Esther Dyson notes: ". . . everybody can get up on the Net, sing their own songs, write their own poetry. You no longer need a publishing house to get a book published. So economics would say that since the supply of content is increasing, the costs of duplication and distribution are diminishing and people have the same amount of time or less, we are all going to pay less" (Dreifus 1996). Of particular concern is the extent to which a balance is struck between intellectual property protection and the needs of users to work effectively in this environment. Some issues that may arise in this context include the following:

1. Will network users be able to browse information contained in digital objects (or that may be manifest when such objects are processed) as easily as readers have been able to peruse books at their favorite bookstore?
2. How will network users be able to "borrow" or otherwise use digital objects stored in repositories? Will there be restrictions on who may access such information?
3. Will authors and other information providers be vulnerable to the loss of significant potential royalties on their works as millions of network users armed with these new capabilities manipulate, disseminate, and interact with their works; or will a digital object infrastructure promote the development of a valuable new market for their information?
4. How will broadcast, cable television, satellite, and other conventional audio and video programming services be integrated in or associated with digital object information services?

### **2.2.2 Legal Issues**

Works and other material configured as digital objects may be produced collaboratively in new and novel ways - for instance, emulating how a motion picture production company handles the various contractual relationships with contributors. Such collaborations may lead to more democratic and effective interaction, e.g., shared learning and discovery, with wider distribution of ideas, and fewer limitations and constraints on communication. Computer networks permit greater and more rapid access to ideas and contributions, and, when combined with new implementations of business rules and practices, may lead to new kinds of businesses and increased employment. The activity raises many questions, however. For example:

1. What constitutes a public performance for copyright purposes in a digital environment? Where digital objects are disseminated over a computer network, should this activity be covered by the copyright right of distribution; and, if so, what should be the scope of this right?
2. How should running a computer program with creative inputs be treated from an intellectual property perspective? Is it a public performance to execute a digital work (such as a video game computer program) that is structured as a digital object?
3. What if such a program is embedded in a compound digital object where its performance is made more popular in its new context?
4. When one contributor provides added value through overlays or the morphing of someone else's intellectual property, at what point in the processing chain does the original work cease to be identifiable, or become sufficiently watered down to bring its link to the final product into question? How will such collaborative efforts relate to current intellectual property law?
5. Should a digital object be treated like a television program for purposes of regulation under communications law? If so, is it necessary to broaden existing communications law concepts of unauthorized interception of a program-carrying signal and the divulgence or publication of its contents for communications law purposes to cover unauthorized access to perform stated operations on digital objects, including repositories structured as digital objects? What stated operations should normally require prior authorization?
6. What constitutes a protected process when providing access to digital objects? What does it mean to communicate a performance of a copyrighted work embodied in a digital object by means of a patented device or process, whereby a bit sequence is received beyond the place from which it was sent?
7. With respect to digital objects, how can we track who owns what and in what contexts? Does "ownership" of digital objects make sense? Would a focus on access to a digital object information service or repository be a flexible starting point in analyzing possible legal implications?
8. How can information owners be adequately compensated when their works are expressed in various digital formats that may be accessed, manipulated, interpreted, and aggregated where such works are configured as digital objects?
9. How should the concept of access to information be applied with regard to confidential and privileged information that has been structured as digital objects?
10. How is denial of service to be addressed? Is this sort of interference an infringement of intellectual property rights? Does this violate communications law or antitrust law? Should denial of service be disallowed and/or protected against?

Several technical solutions and new business models may be introduced that could enable some or all of these activities without sacrificing the interests of individual contributors or dampening their motivation or enthusiasm for sharing their knowledge with society.

### **2.2.3 Business Issues**

Another set of issues arises when planning and implementing business models to develop packages of information in various digital formats. New forms of business will need to evolve in the networked digital world to support the use of digital objects and other digital resources. Some issues to be addressed in this context include the following:

1. How should digital objects and other similar digital resources be identified?
2. Will automated licensing mechanisms be developed within a network environment to facilitate access to digital objects and their contents?
3. What will be the range of services that repositories of digital objects may provide, and what are acceptable rules of procedure and other terms and conditions for accessing such repositories? Will third-party services such as indexing and archiving services arise to facilitate access to these repositories?
4. Will the information contained in a digital object influence the rules for accessing it in a repository; and, if so, how will this be handled in practice?
5. What will be the liability of a certifying authority when authenticating information resources structured as digital objects? In the event that different repositories are subject to different regulatory environments, what impact will this have on security arrangements?

### **2.2.4 Technology Issues**

Many aspects of technology could be selected for discussion here, but we focus instead only on those issues relating to agent technology. This technology represents one of the newest and, in many ways, most interesting and controversial areas of technology. Specific questions that arise in connection with intelligent agent technology include the following:

- How does the use of agents structured as digital objects that operate on information affect intellectual property? Is the output from such an agent a new work? Where is the boundary, if any, between the old work and the new work?
- Should such agents have rights? Under what circumstances may they negotiate licenses on behalf of users?
- Are the agents themselves to be viewed as inventions with their own patent protection, trademarks, etc.? Are they also subject to copyright protection? How does, or should, communications law regulate their behavior?
- How should highly intelligent agents, sometimes called knowledge-based systems, be treated from a legal and business perspective when they give

advice that others resell? What are the issues associated with multiple nestings?

- How should the transformation of one digital object into another be viewed from a rights perspective?

### 3.0 Key Infrastructure Requirements

---

Commercial rights management technologies typically require one or more infrastructure services. Usually, these services involve repository management, data processing, or similar capabilities. Some entail the provision of gateways between the rights management technologies and various general infrastructure services such as financial clearinghouses. Following are descriptions of key infrastructure components of an open architecture that supports digital objects.

- **Persistent unique identifiers** - Digital objects and other digital resources need unique identifiers that can potentially last indefinitely. In fact, a key characteristic of a digital object is the presence of a unique persistent identifier in its metadata. Multiple identification schemes may be desirable in certain circumstances; however, there should be some widely understood methods for resolving these identifiers to permit access to information regardless of the source of the identification scheme. This is particularly important where digital object technologies used in rights management need to interact and collaborate.
- **Global resolution system** - A widely recognized global resolution system for identifiers will allow service providers to identify digital objects - even though based on competing technologies - as unique entities in information commerce. For example, a digital object protected using one technical approach may be located and retrieved by a software agent configured as a digital object using a different container technology. In this instance, the retrieved digital object may retain its original structure and unique identifier when incorporated in the second container for delivery to a customer.
- **Metadata standards** - Digital objects have associated metadata (such as "handles" that uniquely identify them) that may contain information regarding usage terms and restrictions, permissible operations, the sources and contributors of the underlying information components, the rights of each source, the kinds of permissions that must be updated, and how to obtain these rights. The metadata may also be used in negotiating special arrangements. For example, if a user wants additional rights beyond those stipulated in the metadata, there could be a link in the metadata to a person or entity identified as authorized to grant rights and permissions in order to negotiate an appropriate license. To ensure widest interpretation, however, metadata must be based on common standards.



- **Certificate authorities (CAs)** - The availability of certificate authorities is extremely important to the full development of efficient information commerce, especially regarding document management and security. Some companies offer to certify individual or organizational identities for purposes of given transactions. Others are willing to be the "root authority" for other certificate authorities. A central unresolved issue concerns possible liabilities a CA might incur. For example, if a CA certifies that a person is a patent attorney and he or she is not, and if that person harms an unsuspecting client, the CA may be held liable.

#### 4.0 Potential Business Opportunities

---

New business models are likely to evolve to meet the needs of digital technology that may prove of great benefit to society. They may foster the spread of ideas, add value to existing information, provide new services, and generate revenue opportunities. While it is important to encourage these new business activities and their possible collateral value to rights-holders, there is a tradeoff between losing new potential business and value, and losing sources of revenue through information piracy. As Stefik (1996) notes, "the dream of universal digital access to high-quality works dangles just beyond reach. Such works are not usually available, because of publishers' concerns that uncompensated copying will infringe and erode their ability to make a living."

To some extent, differing values are at issue here in determining the so-called public good - free speech, free exchange of ideas, profitability, cultural preservation, equal access, community values, to name but a few. There are bound to be conflicting values in all this, including different community values; technology therefore should be able to support not just one set of values, but be flexible enough to support different values in different circumstances.

Computer networking technology has the potential to provide all manner of new services, many not yet even imagined. These services may not intentionally violate the letter - or even the spirit - of intellectual property or other laws, but they may be perceived as doing so, leaving potential service providers with sufficient doubt or fear of liability that they will either not offer the service at all, or price it too dearly for it to be of much interest. Criteria should be developed that would permit a large class of operations to be performed on digital objects and other digital resources without prior authorization. Similarly, other operations would be considered as requiring prior authorization under most, if not all, circumstances.

Many different pricing schemes may be implemented depending on the nature and scope of rights involved in any given context. A business model might rely on pricing for a dynamic digital object that is constantly updated and refreshed, and that can traverse myriad dynamically updated communications pathways, where each use generally will traverse a new path, often with new or updated information. Other business models might

offer digital objects for free as an inducement for a customer to purchase hard copies (books, etc.) or to sign up for a more comprehensive service. Such digital objects may be regulated much like television programs today under communications law as well as, where appropriate, under patent and copyright laws. Focusing on the implementation of a digital object infrastructure from a communications law perspective may facilitate the evolution of rights management systems for any incorporated contents (Dunstan and Lyons 1994).

In this context, it may be helpful to have in mind an example of a type of business offering that could be provided today. Current technology enables vendors to provide some or all of the following services, several of which are now under development (Bock 1996; IBM infoMarket 1995; and Sibert et al. 1995):

- linking content providers to those who want content;
- providing content or content-related services;
- acting as a repository for digital objects;
- providing abstracts and indices;
- searching content;
- employing encryption and related techniques to manage rights and interests and to ensure the integrity of digital objects and their contents;
- delivering information on disks or CD-ROMs, or providing network access via e-mail, browsers, etc.;
- keeping information protected until the digital object is opened (e.g., in order to open an object, the user must contact a clearinghouse to handle the payment); and
- operating somewhat like a bookstore (e.g., understanding content, generating abstracts, and selling digital objects to the public).

## **5.0 Access to Digital Objects**

---

To learn about a digital object's contents, or to interact with an object to obtain some service, the object's data must be processed. In this regard, "processing" refers to those operations that manipulate content and those that only act on the container. The latter may be considered as "content-free operations." Simple actions such as rendering - whereby a digital object is interpreted to manifest its contents - or identification - whereby a digital object is interrogated to determine its unique identifier rather than its contents - would both normally constitute processing. In the case of a simple digital object (i.e., one that does not contain other digital objects), however, only the latter action would typically be a content-free operation.

More complex actions might include those that deal with multiple digital objects such as those that access content to transform one object into another; or those that merely aggregate multiple digital objects into composite structures, but that do not actually

access their contents. Again, both of these would normally constitute processing, but only the latter would typically be a content-free operation.

Many types of operations can be performed on digital objects. One way to categorize these operations is with an a priori listing of the various types. Another way is via a computer-interpretable language or set of languages that can be used to specify the operations or their types. A third way is via computer programs possibly contained within the digital objects themselves that can become active agents in negotiating rights and permissions for the objects and/or their underlying contents at various locations in a network environment. By delineating specific types of operations that may be performed on digital objects - that is, **stated operations** - a basis may emerge for orderly management of rights associated with digital objects and their contents. These stated operations may dictate the terms and conditions under which digital objects may be stored, accessed, manipulated, communicated, and otherwise shared.

In drawing up any such list, or in delineating types of operations more generally, care should be taken to avoid undue specificity at this early stage so as not to impede the development and potential of this new capability. Further, a simple listing of types of operations does not rule out more complex subsets of each type, as well as the composition of various types in interesting new ways. Through careful deliberation and realistic experimentation, those categories of stated operations associated with digital objects that require prior authorization should be distinguished from those that do not. Many operations on digital objects - typically those of a commercial nature - appear to have overt effects on rights-holders. For example, operations on a digital object may be subject to specific terms and conditions in licensing agreements set forth in an object's metadata or elsewhere, or, simply, a requirement that appropriate attribution of authorship be given and that the integrity of the material be respected.

A consensus may emerge on the types of operations that should be permitted outright, or the conditions under which they would be permissible without prior authorization may be delineated. Efforts in this direction appear advisable. A few examples of categories of operations that may be performed on digital objects are given below. This list is necessarily incomplete, and its further development could benefit from additional study and experimentation.

- **Processing** - Normally, processing of a digital object so as to manifest its contents - i.e., to interpret a digital object for the purpose of using the underlying information in another system or communicating it by some means either directly or indirectly to another person - would not be considered a permissible operation without express authorization. It would also be impermissible if processing entailed deleting or otherwise rendering unintelligible the terms and conditions or other information associated with a given digital object; or deleting or destroying a stored digital object from a repository without authorization.
- **Distribution** - Distribution of a restricted digital object without authorization would usually not be considered a permitted operation. On

the other hand, denying access to a digital object where there are no restrictions placed on operations that may be performed on the object and/or its underlying contents, e.g., inhibition of service via methods such as jamming servers or communications pathways, should be proscribed.

- **Replication** - This operation refers to the replication of digital objects for ease of use and/or reliability. Replication is often a critical system function that adds value and may not necessarily involve intentional violation of intellectual property rights.
- **Compression** - Many compression schemes, both lossless and lossy, are based on content. While the use of compression may offer value in some circumstances (such as lossless compression), it does not necessarily result in an infringement of intellectual property rights. Indeed, in many cases, compression can actually increase the value of intellectual property by enhancing the ease, timeliness, and cost effectiveness of its distribution.
- **Packaging** - Many digital information services are not intended to access the contents of digital objects but merely to assist in packaging or repackaging them. Packaging techniques might include adding formatting information, encrypting information, moving information from one container to another, or reordering discrete digital objects contained within a given digital object. In many cases, there may be no natural order specified for the multiple components of a digital object, and the specific packaging choice may be at the option of the sender or the receiver (or both), provided that the contents are not otherwise changed.
- **Caching** - Another useful service is caching the information of others for local redistribution or sale. This may be especially appropriate for certain classes of digital objects, such as those that contain - or may be interpreted to manifest - information considered as digital money or registered bonds, where the object can only be transferred from one repository to another without alteration, and where only one original is deemed to exist logically in the system.
- **Carriage** - Intermediate carriers that provide point-to-point delivery based on the wishes of the originator and/or subscriber should be able to treat that operation much like traditional common carriage. Resale carriers, however, might have to be specifically recognized as such by law or regulation.
- **Aggregation and integration** - Aggregating and integrating streams of information coming from different sources provide too much value to be prohibited entirely, even though such operations make it more difficult to enforce rights. An example of such aggregation/ integration is combining weather prediction data structured as one logical entity with another group's digital object embodying oil storage and distribution plans, and presenting the resulting digital material in a structured form over a geographical information system.
- **Clearinghouse services** - Certain activities of rights and financial clearinghouses may be candidates for exemption from liability under relevant laws. Rights clearinghouses, for example, perform several useful

intermediary functions, including encouraging rights-holders to put digital objects into repositories and attracting a larger potential customer base. As for financial clearinghouses, commerce in information "goods" or "services" - like commerce in anything else - requires secure, efficient, timely, and accurate clearing of financial transactions. Rights-holders and other value chain participants want assurance of payment or receipt, but typically do not want to manage a large number of financial interfaces with widely dispersed customers. At the same time, users want to be able to pay for information goods or services - in a variety of ways - via a common, trusted interface. The parties to a transaction also need to be able to verify that a given exchange has occurred as it was mutually intended, and to preclude repudiation of the transaction by either party.

- **Reference services** - The explosive growth and increased accessibility of information, ideas, and concepts creates a demand for more and better indices, catalogs, and contexts in which they can be placed. In a networked digital world, technology permits myriad indexing and cataloging schemes and contexts to be developed quickly and efficiently, as well as linked to a dynamically updated worldwide information mesh of digital objects and other digital resources. Libraries and similar information providers may furnish a variety of future information services, e.g., locating intelligence in the network; launching agents to perform research tasks; translating in multiple contexts; and providing security, authenticity and the building of network environments from disparate resources tailored to user needs. Some of these operations may be permitted without prior authorization. For example, cataloging and indexing digital objects for the purpose of delineating the collection for subsequent access may be permitted, while performing such actions for the purpose of describing the contents of a collection in substantial detail may not.
- **Brokerage services** - Organizations may wish to act as brokers for digital objects owned by others without authorization, provided that they do not make available the digital objects themselves. An example of such an offering might be the generation of digital objects called **meta-objects**, whose primary purpose is to provide references to other digital objects. Brokerage activities may require the use of minimal amounts of information from the digital objects, such as their names, titles, etc. Such services may be restricted.
- **Maintenance** - These services may require access to repositories of information for diagnostic or repair purposes. Any legitimate access to digital objects for such purposes, or for improving system availability, would likely be considered a permissible operation without authorization.
- **Authentication** - Authentication services may require access to the contents of digital objects for the purpose of watermarking, certification, and/or time stamping. Whether, and under what circumstances, such services are exempt from liability for their operations is an important area for discussion.

- **Transformation/browsing** - Transforming and browsing digital objects are two of several different kinds of processing actions that merit further discussion regarding permissibility. For example, transforming one digital object into another (i.e., morphing) may leave no perceptible trace of the former but still not yield an exact replica of the latter. To what extent is there a remnant of the former resident in the result? A related example involves the stripping out of selected portions of a digital object during some intermediate operation such as morphing. And what about searching and/or browsing some portion of a collection of digital objects for the purpose of locating a specific object or identifying material of interest - should this be considered similar to searching or browsing in a bookstore, where the full content is not actually digested?

Various other operations come to mind in a network environment such as reporting (via agents), traffic analysis or forecasting services that can reconcile various operations previously taken on digital objects or those involving specifications for billing and payment, and accessing the metadata of a digital object to determine the terms and conditions governing its contents.

## 6.0 Reflections on Rights Management Technologies

---

What should a rights-holder (reasonably) be required to do in order to protect his/her rights in a digital environment? If terms and conditions that cannot be removed or modified can be easily incorporated within a digital object's metadata, would a user be justified in presuming that the only constraints on access are those found in the metadata, and contact the owner only if he or she wants different terms and conditions? Should independent parties who add value to original works be permitted/required to add their own terms and conditions to the business rules and procedures in metadata as the modified/augmented works are passed down the value chain?

The most technically advanced rights management systems will most likely be delivered in a powerful, flexible, and efficiently protected manner which supports a digital object infrastructure that allows appliances and devices of all kinds large and small to participate in information commerce. A distributed rights management system is that collection of technologies and processes that can assist in determining and enforcing rights and interests and in ensuring the persistence and integrity of information. It may be comprised of a single system whose components are distributed, or a collection of systems that have well-defined and open interface standards. A digital object may itself be an active component in a rights management system, carrying along with it the terms and conditions for its access or enabling dynamic negotiation of rights and permissions. A list of some general business considerations and infrastructure requirements for distributed rights management appears in appendix A.

A range of technologies are already deployed or will be in the marketplace shortly for digital rights management in the network environment. Several noteworthy examples of these technical capabilities are described in appendix B. In one way or another, they all depend on protecting digital information mapped into analog signals (i.e., continuous waveforms) at different levels of granularity. Communication of digital information from storage media or over networks (and other pathways) requires detection of bits from analog signals. Further, in a digital world, bits may be encrypted in a wide variety of ways for storage or communication. Inherently, this distinction enables structured digital information to be distinguished as a logical entity from the waveform for purposes of rights management.

Rights management systems may make use of certain current and future enabling technologies - technologies that do not, as such, manage rights. Various technical approaches for rights management have been deployed; others are now under development. Two important technical approaches that should be addressed in a rights management context are discussed below.

### **6.1 Securing the "Pipe"**

A basic rights management technique focuses on protecting the entire set of signals transmitted over a communication pathway between a user and a server. This, in effect, protects the entire interaction with an information source. It may be seen as directly protecting the "pipe" rather than the initial, ultimate, or any intermediate container or underlying content. Typically, an entire protected interaction (after some preliminary exchanges) between an information provider and a customer would be encrypted or scrambled.

In an open environment such as the Internet, someone intercepting an encrypted signal may not easily access the unencrypted bit sequence, much less any incorporated digital objects or underlying content. This technique ensures a certain degree of confidentiality, so that unauthorized parties cannot necessarily determine any material aspect of the information interchange, such as which content was being provided. It also allows for a degree of security when passing credit card and other financial information. Once the signal is received by the customer, however, it will generally be decrypted and its underlying structure made available to the end user. Some of these structures, such as individual digital objects, may be in the clear or they may be separately encrypted.

### **6.2 Protected Containers**

A protected container approach, on the other hand, emphasizes the individual package that incorporates information that may itself be encrypted or embedded in another container. A simple example would be the lock/unlock approach used by a publisher or other distributor that encrypts specific content for delivery - in encrypted form - to a customer. Lock/unlock systems typically include a financial clearinghouse service for processing payments; these may also report some usage information to rights-holders. Once the customer has paid for the digital information service, he or she can decrypt the content (by first getting a key or password).

Most encryption systems do not provide for ongoing control, metering, and billing capabilities; others provide these features by requiring the customer to run a dedicated application each time he or she uses a digital property. Thus, the information service is often not widely available for general use by typical PC applications. These services usually ensure that rights-holders are paid upon delivery, and sometimes on a pay-per-use basis.

A more sophisticated approach for providing protection is based on the idea that it is possible to create a data structure ("digital object") that allows portions of the structure to be identified and separately protected. As a general matter, any digital object may be protected through the use of encryption. A variety of different symmetric and asymmetric encryption systems can be employed, with each having a distinct role in protecting a digital object and/or its contents - or, in some cases, the rights management information delivered with the content. Such encryption capabilities enable a more complex rights management approach, and raise the threshold for those who would attack a container. A container may hold unencrypted fields for titles, abstracts, thumbnails (abbreviated or low-resolution images), or any other information the provider may wish to present to the customer or potential customer. An application on the user's machine opens the container and makes the content available to the customer, sometimes only within the context of that application, or more generally (e.g., an Internet browser or agent software).

Some protected container systems are "clearinghouse-centric": they require a conversation with some central server in order to purchase and use the content. This model reflects a client-server orientation in which large central servers provide most of the important functionality to less capable client software. Rights management capabilities are typically limited to securing distribution and ensuring that rights-holders are paid for the use of information. In this scheme, a customer may be permitted to further distribute the protected container and its contents to any number of people, each of whom may contact the information distributor's clearinghouse to effect payment and purchase. Typically, the customer cannot modify any of the controls associated with the content, nor can permissions and prices not originally distributed with the content be efficiently obtained.

Other protected container technologies go beyond server-centric rights management models. These allow each enterprise in a value chain to contribute business rules under the control of more senior participants, thus enabling flexible implementation of the broadest range of business models. By enabling the creation of chains of title, these advanced rights and information commerce technologies enable both traditional business models and new, as-yet-undreamed-of models, to emerge in cyberspace (Bock 1996; and Sibert et al. 1995). These systems support the notion that consumers can act not only as peer-to-peer distributors, but that, if authorized, they can contribute their own business rules and procedures. This capability allows consumers of business, educational, entertainment, and other information resources to become value-added resellers.

A brief example demonstrates the point. A scientist at a for-profit research company writes a long review article, and packages it in a secure container along with a licensing



agreement that (1) restricts others from modifying the content, and (2) sets a one-time charge of \$1 to read or print the article. Having obtained any necessary permissions, the scientist also incorporates two related articles in the same digital object. Each of these articles could carry a similar set of conditions regarding modifications and pricing. Readers of the first article are not required to access the two additional articles. However, for the convenience of providing the additional articles, the researcher marks up each by 10 cents. So, if someone opens this container and views all three articles, the researcher would receive \$1.20 and each of the authors of the other two articles would be paid \$1.

## 7.0 Conclusion and Recommendations

---

This paper addresses several important issues about managing access to digital information. It highlights the overlap of legal and technical concerns, and introduces certain important concepts that may facilitate the development and evolution of network-based information services, in particular, rights management systems. One such concept is that of a **digital object**, i.e., a data structure for identifying and organizing information for access over a communication network. Another concept described here is **stated operations**, that is, the types of operations that may be performed on digital objects. Stated operations are a useful construct in helping value-added service providers and other users of protected information understand the scope of their liabilities.

Key infrastructure requirements in pursuing new business opportunities using network-based digital information are also discussed, and various rights management technologies introduced. It is highly advisable to gain experience with these capabilities so that their implications may be better understood and the resulting systems and techniques refined. Some guiding principles for the evolution of the field are introduced. Of these, the most important are the focus on digital objects as packages of structured digital information (encrypted or not) for provision of information in the network environment and the use of rights management techniques such as the incorporation of terms and conditions in a digital object's metadata, together with a capability of negotiating additional permissions with a rights-holder, or a means of determining where such information may be accessed.

The new capabilities described in this paper may require a new or revised legislative framework. Before laws are revised or enacted, or new regulations issued, however, additional experimentation and experience are needed with digital objects in the marketplace. This will let nascent markets develop without undue regulation, provide a base of experience to guide the formulation of new laws, and suggest where change is - or is not - justified. In particular, and as noted earlier, overlaps in existing laws, such as copyright, patent, and communications law, should be exploited to facilitate a more contoured and textured approach in dealing with issues of protection and liability than is possible through reliance on any single body of law.

A guiding principle in interpreting existing laws and regulations - and in formulating any new provisions - should be as follows:

If an operation does not restrict or impair a rights-holder's revenue stream or existing rights or interests, is not explicitly restricted by the rights-holder in the stated terms and conditions of use that are linked to or included within a digital object, or is not otherwise explicitly restricted by law, then the operation may be presumed to be permitted. This principle should apply independently of the technology used to realize or convey a digital object, or the means used to transact business.

Public policy should foster and support an environment that promotes experimentation, enhances understanding, and encourages the development of pilot projects and new business practices. And, for the duration of such efforts, limited immunity under antitrust law and certain intellectual property laws may be desirable. For their part, the projects thus supported by public policy and legal immunity should be aimed at developing and deploying new technology and infrastructure to facilitate the conduct of business in a digital environment, including, in particular, unique identifiers and resolution systems for digital objects, protocols for access to digital objects, and public key cryptography as well as certification and authentication infrastructure that meets the legal requirements for U.S. domestic commerce. Public policy should also promote the development and voluntary use of an open architecture, including open standards and common business practices to support these efforts.

## References

---

- Bock, G.E. 1996. "InterTrust Commerce Architecture and Developer's Kit." In *Distributed Computing Monitor*, p. 3. Boston: Patricia Seybold Group.
- Dreifus, C. 1996. "The Cyber-Maxims of Esther Dyson." *The New York Times*, July 7: 16-18.
- Dunstan, J.E., and P. Lyons. 1994. "Access to Digital Objects: A Communications Law Perspective." In *1994 Annual Survey of American Law*, pp. 363-82. New York University School of Law.
- IBM infoMarket. 1995. "IBM Cryptolope Containers." <http://www.infomarket.ibm.com>.
- Kahn, R.E., and R. Wilensky. 1995. "A Framework for Distributed Digital Object Services." <http://www.cnri.reston.va.us/home/cstr/arch/k-w.html>. Reston, VA: Corporation for National Research Initiatives (CNRI).
- Sibert, O., D. Bernstein, and D. Van Wie. 1995. "The DigiBox: A Self-Protecting Container for Information Commerce." *Proceedings of the First USENIX Workshop on Electronic Commerce*, p. 171. New York, NY: USENIX Association.
- Stefik, M. 1996. "Letting Loose the Light." *Internet Dreams*, p. 221. Cambridge, MA: MIT Press.

## Appendix A

---

Following is a short list of some of the most desirable attributes of a rights management system from a business perspective - that is, capabilities that facilitate protected information commerce and rights management.

- **Associated terms and conditions** - A key idea is that terms and conditions can be associated with specific digital objects, as well as with their underlying contents and that a rights management system may assist in enforcing these provisions. Among the terms and conditions applying at the digital object level are the "stated operations" described earlier. Permissions for access to a digital object or its contents may be included within a digital object's metadata, or may be delivered independently. The most sophisticated rights management systems will allow rights-holders to describe access terms and conditions for different models and contexts. For example, a property could be delivered with two sets of rules: one for pay-per-use and another that specifies a larger one-time fee for unlimited use. A property could also be associated with rule sets for different contexts, so that different pricing models or permitted operations could be granted based on whether the user were, for example, a student or a government employee.
- **Assured integrity** - Rights management systems also need to protect digital properties against loss of integrity through alteration or substitution of one work for another. Consumers may need assurance that the work they purchase or rent has not been altered; and rights-holders need assurance that their brands will not be undermined by unauthorized modifications or fraudulent recreations of their properties.
- **Chain of operations and value management** - Flexible rights management systems will support the ability of each distinct entity in the value chain - author, publisher, aggregator, repackager, payment method provider, customer, etc. - to contribute independently the business rules and other conditions that reflect its individual rights and interests. Some of these contributions may be under the control of more senior participants in the chain. For example, authors may permit their publishers to modify their works, but not permit others in the distribution chain to do so. Publishers may offer discounts to aggregators who, in turn, can mark up the price of the work.
- **Efficient and tamper-resistant** - Rights management technologies vary in their defenses against sustained attacks. Of course, not all digital information has to be made completely secure. Rather, the protection has to be high enough that the cost of breaking the

security is disproportional to the value of the property being protected.

- **Flexibility** - Business, educational, entertainment, and other types of information are "consumed" using a variety of devices and appliances large and small. While some material is made available over a computer network, a variety of information resources are delivered on physical media - for example, motion pictures may be delivered on digital video disk. Rights management systems must incorporate technologies that support devices of varying sizes and capabilities as well as huge numbers of transactions. Further, an advanced rights management system must be able to support the broadest possible range of business models on an enterprise-by-enterprise, category-by-category, or even property-by-property basis. Such models refer not only to pricing strategies, such as "one-time purchase" or "pay per use," but also to the relationships among value chain participants. Value chain participants should readily be able to create ad hoc business and rights relationships, and dynamically change them in response to market or business conditions.
- **Payment methods** - Rights management systems must ensure that rights-holders and other distribution chain participants are paid appropriately for those properties that carry a price. In doing so, however, rights management systems should not impose undue economic burdens on the commerce being protected. For example, the cost of providing infrastructure services such as financial clearinghouse services must not be disproportionate to the value of the commerce being transacted. Many payment methods now exist, and more are likely to emerge. The best rights management systems should provide users, rights-holders, and others in the distribution chain with the broadest range of cost-efficient payment methods.
- **Persistent protection of digital information** - Some of the earliest rights management solutions delivered digital information in encrypted form, but made the unencrypted version available in unprotected form after payment. Advanced rights management technologies should be able to provide rights-holders with more persistent protection, if that is what they want.
- **Security and trust** - An obvious, but essential, requirement of a rights management solution is that it be trustworthy and secure. And a trusted system is only as secure as its weakest link. While existing security problems on the Internet - and on intranets - are well-known, much of the current focus has been on securing the "pipe." In contrast, advanced rights management systems should be capable of protecting both the digital object and/or its contents. Once encrypted (e.g., with protected keys), the sealed package as a whole (or discrete portions thereof) is protected. Rights-holders

may want to enable some operations on their properties but to forbid others. As rights management systems become fully integrated with standard applications, controls will become increasingly fine-grained, thus providing rights-holders with many options.

- **Support for advertising models** - It is envisioned that there will be many different models of repositories, e.g., serving as a bank, as an advertiser, as a video distributor, or as a broadcaster. Many business models for digital information commerce will depend on advertising for some or all of their revenue. Thus, rights management technologies should support many advertising models.

## Appendix B

---

Several basic technologies are now available that could be used successfully in advanced rights management systems. Some of the most promising are described below.

### Encryption

Encryption is primarily used to make information secret or confidential so that only those who possess a certain key can access it. This information might be a digital work, rights management controls or pricing information, or other element. The key is simply a number. Advanced rights management systems make use of two different kinds of encryption technology: symmetric, or secret key, and asymmetric, or public key. A **symmetric key system** uses the same key to both encrypt and decrypt information; consequently, it is important to keep this key protected. An **asymmetric key system** uses a pair of keys that share certain mathematical properties. In an asymmetric system, if information is encrypted with one of the keys, it can only be decrypted by the other. As in symmetric key systems, one of keys - the private key - is usually protected. The other key - the public key - can be made available publicly without compromising system security, since the private key's value cannot be determined from that of the public key. Anyone can use the public key to encrypt information, but only the holder of the private key can decrypt it.

### Digital Signatures

Digital signatures provide very strong indicators of integrity and authorship, and thus can be very important to rights-holders, value chain participants, and customers. For example, stockholders may need to know that a financial report has not been modified accidentally or intentionally. Digital signatures make that knowledge possible; they also can provide a reliable way of ensuring that rights are exchanged based on enforceable licensing agreements.

A digital signature can be generated using the cryptographic method called the **one-way hash function**. The hash is a calculated number that reflects the content of a particular digital object. If even one bit is changed, the value generated by the one-way hash function will also change. And, conversely, it is extremely difficult to change a digital object in such a way that corresponds to a particular hash value. These features allow a person to "sign" a particular digital object using his or her private key. First, the person who wants to sign an object calculates the hash value using a one-way hash function, and then encrypts this hash value using the private key. Anyone who has the signer's public key can then decrypt the value, recalculate it, and compare the two numbers. If they are the same, the recipient knows that the digital object is unchanged and that it could only have come from the person with the corresponding private key. This protected hash method is even stronger than the one-way hash since it involves encryption.

## **Certification**

Asymmetric key systems can also be used in conjunction with "certificates" that attest to or warrant some fact about the owner of a public/private key pair. These certificates are issued by a "certificate authority" (CA), a service that performs varying degrees of fact checking before issuing a certificate. Among the kind of facts that can be certified are individual identities or membership in a particular class or organization. Certificates can also attest to the authenticity of a digital object and/or its content. For example, a CA may create a message that contains a structured factual statement that the public key in a given message belongs to John Q. Public, or that the person using a specific public key was born on April 19, 1960. The CA then encrypts the information with its private key. As with the digital signatures described above, this procedure allows those people who have the CA's public key to decrypt the message, and - assuming they trust the authority - to accept the facts conveyed as true.

In the best rights management systems, certificates can also convey contexts for rights usage. For example, a publisher may wish to offer discounted or free goods or services if customers can provide a certificate attesting to the fact that they are affiliated with an institution of higher education, or a certificate indicating that they are affiliated with a not-for-profit research organization. Similarly, the rights management system may enforce one set of conditions on prices, taxes, and currencies in the United States, and another set in France, depending on which certificate covering jurisdictional issues is provided by a customer.

## **Fingerprinting and Watermarking**

Steganography is the science of hiding messages in communications and is the basis of various so-called fingerprinting and watermarking techniques. For example, the location of marks on a page can be adjusted in minute ways so as to encode hidden information. Certain pixels of a digital image can be manipulated for the same purpose. Such hidden messages may be difficult to detect and difficult to remove by anyone who does not know how they were applied in the first place.

Advanced rights management systems can make use of fingerprinting in several ways. First, these methods can be used to encode copyright information such as property title and other identifying information (e.g., who owns the copyright and the year of first public availability). Second, when content is released out of a controlled environment - usually for a fee - a fingerprint can be added to many kinds of properties indicating who exported the property and when. The fingerprint can be used to indicate, or at least suggest, the initial source of the information in the event of infringement.

## **Other Security Measures**

Software-based rights management systems can protect against overall system defeat and unauthorized access to digital information. The goal is to ensure that the cost of defeating the system is disproportionately larger than the value of the properties protected by the system. Some rights management systems are able to make use of secure hardware when it is present. This hardware may be a microprocessor and memory on a PCMCIA card, a tamper-resistant co-processor on the motherboard, or recently introduced advanced



microprocessors. Secure processors help ensure that the digital properties will continue to be protected from all but the most dedicated and sophisticated attacks. More importantly, secure processors can provide strong protection for the rights management information that keeps the system operating correctly. This protection allows more processing of rights-related information on users' systems, which lightens the load on central servers and opens up new opportunities for user-initiated commerce. Until secure silicon is widely available, and perhaps beyond, software-only solutions can do a good job of protecting a broad range of digital information.